

AAS 98-062



**AUTONOMOUS FORMATION FLYER (AFF)
SENSOR TECHNOLOGY DEVELOPMENT**

**George Purcell, Da Kuang, Stephen Lichten,
Sien-Chong Wu, and Larry Young**

**Jet Propulsion Laboratory, California Institute of Technology
Pasadena, California**

21st ANNUAL AAS GUIDANCE AND CONTROL CONFERENCE

February 4-8, 1998
Breckenridge, Colorado

Sponsored by
Rocky Mountain Section



AAS Publications Office, P.O. Box 28130 - San Diego, California 92198

AUTONOMOUS FORMATION FLYER (AFF) SENSOR TECHNOLOGY DEVELOPMENT

**George Purcell, Da Kuang, Stephen Lichten, Sien-Chong Wu, and
Larry Young***

The DS-3 mission will demonstrate various elements of the technology required for space interferometry, including Autonomous Formation Flying (AFF). It will consist of three spacecraft, each having a degree of autonomy but all comprising a single instrument and constrained to move together at the vertices of an equilateral triangle with sides of 100 to 1000 m. In order to meet the mission's goals, AFF must maintain the distances between spacecraft within 1 cm and the relative orientations of the spacecraft within 1 arcminute per axis.

This paper proposes an implementation of AFF that borrows technology from the Global Positioning System (GPS), using measurements of both r-f carrier phase and a ranging code. Each spacecraft will have at least one transmitting antenna and three receiving antennas, operating at 30 GHz with a code rate of 100 Mcips/second.

To validate the proposed scheme, two sets of covariance analyses have been performed, using batch and sequential processing. Both analyses suppose measurement errors of 1 cm on the ranges and 10 μ m on the phases, but they process the data differently and make somewhat different assumptions about the nature of systematic errors. Nevertheless, the two analyses reach similar results, concluding that with careful calibration the mission requirements can be met by measurements of the specified accuracy.

The paper also summarizes various aspects of the proposed implementation that have not been entirely worked out, indicating in general terms how each might be handled and what problems remain to be solved. These problems include initialization, the location and beam shape of the antennas, multipath, systematic measurement errors, and the effect of the local transmitter on received signals.

BACKGROUND

DS-3 is a proposed \$100M-class mission within the New Millennium Program, whose purpose is to demonstrate various elements of technology required for space interferometry. One of those elements is Autonomous Formation Flying (AFF), the subject of this report.

* All at the Jet Propulsion Laboratory, California Institute of Technology, M/S 238-600, 4800 Oak Grove Dr., Pasadena, California 91109. Phone 818-354-3395, FAX: 818-393-4965, e-mail George.Purcell@JPL.NASA.Gov.

Introduction to DS-3

As currently conceived, the DS-3 mission will consist of three spacecraft, each having a degree of autonomy but all comprising a single instrument and constrained to move together in a tightly controlled formation. Two of the spacecraft will be the collecting elements of the interferometer, roughly cubic in shape and having masses of about 150 kg. Each will use a mirror with a diameter of 12 cm to reflect the collected light (wavelengths 500 to 900 nm) to a third spacecraft, the combiner, of comparable size but somewhat heavier (about 250 kg). Because of the modest collecting area, the faintest measurable sources will have visual magnitudes in the range 10 to 12. The interferometric baselines will vary in length from perhaps 100 to 1000 m, and the three-spacecraft array will form an equilateral triangle. During a planned lifetime of 6 months, the instrument will demonstrate its ability to point at specified targets, change baseline length, and maintain the formation at the required accuracy, as well as to find and track the interferometric fringes and report its measurements back to Earth. In the process, the instrument will measure the correlation amplitudes of 50 to 100 objects, mostly stars in our own galaxy. These amplitudes, in turn, will indicate the size and structure of the objects.

Clearly, monitoring and controlling the array configuration will be a crucial element in the operation of the DS-3 interferometer. That is, the collector spacecraft must be separated by the specified distance in the specified orientation and pointed toward the designated source in such a way that the two light paths reach the correct points at the combiner spacecraft. Furthermore, the total geometric lengths of the two light paths must be kept so nearly equal (within about 30 cm) that the optical path lengths can be equalized exactly by the optical delay line at the combiner. Incidentally, the need to maintain the array configuration without continual thruster firings mandates that the array operate away from strong gravity gradients—that is, beyond Earth orbit. As a result, DS-3 will operate in a solar orbit similar to the Earth's but trailing Earth by something like 0.1 AU. Thus there are two types of knowledge requirements on the configuration. First are the requirements on the orientation of the array in inertial space. These requirements are met by star trackers on at least one of the spacecraft. Second are the requirements on the internal geometry of the array, including the lengths of all three sides of the triangle and the orientations of the three spacecraft with respect to one another. At and below the millimeter level, many of these requirements can be met partially by optical means; but at the centimeter level and above optical tracking is impractical, and this is the regime of AFF.

Introduction to AFF

Conceptually, Autonomous Formation Flying is a process in which an array of spacecraft makes continuous measurements of its "array configuration" and uses those measurements either to maintain an existing configuration or to move smoothly to a new one, all without external measurement or control. Generally speaking the array configuration would include both the distances between all pairs of spacecraft in the array and also the orientations of the spacecraft in a coordinate frame defined by the array's internal geometry. Thus the term "formation flying" in its usual connotation is entirely apropos, and the formation's "autonomy" resides in the lack of external control. In the case of DS-3, for example, ground controllers might—at least in principle—simply send the array an observing schedule and then sit back to wait for the results.

Application of AFF requires three distinct steps. First, perform measurements sufficient to determine the required array parameters. These parameters include the coordinates of the spacecraft in an internally defined reference frame, three orientation angles for each spacecraft, and possibly some "nuisance" parameters required by the estimation scheme. The measurements include distances between various points on different spacecraft and probably other kinds of observables such as radio-frequency phases and phase rates.

The second step is to apply an estimator that accepts the measurements as input and generates estimates of the parameters as output. This step, using well known algorithms, is rather straightforward. It is also possible for the estimator to incorporate additional measurements from outside the AFF machinery to strengthen its solution.

The third step in implementing AFF is to use a time series of estimates of the relevant array parameters to control those same parameters: either to maintain an existing configuration within specified limits or to shift smoothly to another configuration. This step will be discussed in a future paper.

In the sections that follow, we discuss a particular way of implementing step one and analyze step two to see whether this approach can attain the accuracy required by DS-3. Although the analysis focuses on a specific scenario (DS-3) with a minimal number of spacecraft (three), the same ideas can readily be extended to larger numbers of spacecraft and applied to a variety of planned JPL missions, including DS-4/Champion and Mars sample and return. In addition, the same principles can be applied to a range of similar control problems, such as formation flying of aircraft.

APPLICATION OF AFF TO DS-3

In this section we consider how AFF might be applied to the DS-3 mission. First we propose a basic approach and calculate estimates of the measurement errors. These calculations will be careful but not rigorous, since they depend on details of the system design that are not yet determined. Then we postulate a more specific design and lay the groundwork for the subsequent covariance analysis and simulated AFF solutions. This design is incomplete and not altogether realistic: it is intended only to provide a framework for verifying that the basic approach is workable.

Requirements, Measurement Scheme, and Measurement Errors

DS-3 requires that AFF provide the distances between spacecraft with an uncertainty of 1 cm (1 standard deviation) and the orientation of each spacecraft with an uncertainty of 1 minute of arc or 2.9×10^{-4} radians (also 1 standard deviation, on each of three rotation axes) with respect to a coordinate frame defined by the relative locations of the spacecraft. The scheme we propose to make these measurements is a GPS-like system in which one or more transmitters on each spacecraft transmit to three or more receiving antennas on each of the other two spacecraft. As in GPS, the transmitters send a carrier modulated by a pseudo-random noise (PRN) code for ranging purposes, and the receivers will use the demodulated carrier to provide a more precise ranging signal with integer-cycle ambiguities. Each transmitter will use a different code so that the receivers can distinguish among the various signals. (It is also possible that the transmitters will use different carrier frequencies.) Unlike GPS, however, DS-3's accuracy requirements demand a shorter carrier wavelength and shorter PRN chip length. In the discussion that follows we will assume a Ka-band carrier wavelength of 1 cm and a PRN chip rate of 100 MHz (that is, a chip length of 3 meters).

With some additional assumptions one can compute the errors expected on range and phase measurements averaged over a given interval. If we consider at first only the effect of system noise and ignore multipath and other systematic errors, the range error for standard GPS-type processing, in meters, is

$$\sigma_R = \pi \frac{\lambda_{chip} d}{\lambda_{rf}} \left(\frac{2kT}{G_T G_R P_T t} \right)^{1/2} \quad (1)$$

where λ_{chip} is the “length” of a chip of the PRN ranging code (m),
 λ_{rf} is the carrier wavelength (m),
 d is the distance (m) between the transmitter and the receiver (assumed to be in the far field),
 k is Boltzmann’s constant, 1.380662×10^{-23} J K⁻¹,
 T is the receiver noise temperature (K),
 G_T is the gain of the transmitting antenna in the relevant direction,
 G_R is the gain of the receiving antenna in the relevant direction,
 P_T is the transmitted signal power, (watts), and
 t is the signal integration time (s).

Similarly, the uncertainty of a carrier phase measurement, also in meters, is given by

$$\sigma_\phi = d \left(\frac{kT}{G_T G_R P_T t} \right)^{1/2}. \quad (2)$$

To get a feel for equations (1) and (2), consider an example in which

$\lambda_{chip} = 2.99792458$ m,
 $\lambda_{rf} = 1$ cm,
 $d = 1$ km (the nominal maximum),
 $T = 100$ K,
 $G_T = G_R = 1$,
 $P_T = 1$ milliwatt, and
 $t = 1$ s.

In that case σ_R is 1.57×10^{-3} m, and σ_ϕ is 1.18×10^{-6} m.

These system noise errors are not realistic because other errors, such as multipath, will dominate. For example, if all three spacecraft transmit and receive simultaneously, some of the transmitted power will inevitably find its way into the receivers on the same spacecraft. Even though signal processing extracts the desired remote signals, the self power will still raise the effective noise temperature of the receiver. In fact, it is not unreasonable to suppose that the self power will be 30 dB stronger than a remote signal. In that case it turns out that σ_R is 6.7×10^{-3} m and σ_ϕ is 5.0×10^{-6} m. These values are more realistic but probably still optimistic. For the purposes of the covariance analysis that follows it will be assumed that the uncertainty of all range measurements is 1 cm and the uncertainty of all phase measurements is 10 μ m.

Geometry, Measurements, and Observables

Figure 1 shows the coordinate frame adopted for modeling the AFF geometry on DS-3. The two collector spacecraft are labeled A and B, and the combiner is C. A reference

point associated with spacecraft A is at the origin of a Cartesian reference frame (x, y, z); B is on the $+x$ axis, C is in the x - y plane at positive y , and the frame is right-handed.




Figure 1. Definition of AFF coordinate frame

In addition to the array frame, each spacecraft has its own local co-rotating Cartesian frame in which the locations of its transmitting and receiving antennas can be specified; and its orientation with respect to the array frame is defined by the rotation angles ψ , θ , and ϕ , as shown in the figure. The origin of each local frame is a reference point on the relevant spacecraft, and when the three angles are all zero, the local and array frames coincide. In general, any orientation can be reached by a rotation through angle ψ around the z axis, followed by a rotation through angle θ around the rotated y axis, followed by a rotation through angle ϕ around the rotated x axis.

All the AFF observables are either ranges or range-like measurements. At any measurement epoch, a range can be measured between any transmitting antenna and any receiving antenna on a different spacecraft. To reduce power consumption and the damaging effects of interference, it is clearly desirable to minimize the number of transmitters; so in this discussion we will suppose there is one transmitting antenna per spacecraft. In that case there must be at least three receiving antennas per spacecraft in order to determine all the rotation parameters. Again, this discussion will suppose the minimum

number. When the array is initializing itself, for example, more transmitting and receiving antennas may be useful; but in ordinary operation it seems likely that the minimum complement of hardware will be optimal.

It is then evident that at any measurement epoch 18 independent range measurements are available. If we denote the transmitters by the names of their spacecraft, using an upper case letter, and the receivers by a lower case letter and a serial number, then the designations of the antennas on spacecraft A, for example, are A, a1, a2, and a3. Using this nomenclature we can then name the six observables measured at spacecraft A: Ba1, Ba2, Ba3, Ca1, Ca2, and Ca3; and the names at the other two spacecraft are exactly analogous.

Clock errors at the transmitting and receiving spacecraft affect the observables, but it should be clear that only differences of clock errors are significant. And although there are three such differences, only two of them are algebraically independent. As a result we need to define two additional parameters of our observation model, namely R_{TAB} and R_{TAC} , which as the symbols suggest are expressed in range units. Other clock terms in the measurement model can be specified in terms of these two. In the analysis that follows, it will be assumed that a single clock offset characterizes all the transmitting and receiving antennas on each spacecraft. With careful design and calibration, this assumption may be adequately justified.

Under the assumptions made above, the array range model is specified at each measurement epoch by 14 parameters: the three spacecraft coordinates x_B , x_C , and y_C ; the nine orientation angles ψ_A , θ_A , ϕ_A , ψ_B , θ_B , ϕ_B , ψ_C , θ_C , and ϕ_C ; and the two clock differences R_{TAB} and R_{TAC} . Since there are 18 observables, a range-only solution can be estimated at each epoch.

Unfortunately, the cm-level accuracy of this solution cannot meet the angular requirement of 1 minute of arc. (Note that a 1-arcminute error on a 50-cm radius vector corresponds to a range error of about 0.15 mm.) Therefore phase observables must be brought into the solution. Undifferenced phase measurements have arbitrary offsets, but astute differencing can make them more tractable. For example, time differences of a continuously tracked phase are unambiguous. Differencing the same transmitter between two receiving antennas on the same spacecraft gives an observable with an offset that is an integral number of carrier cycles, if all the instrumental delays in the two signal paths are identical. Higher order differences are also possible. The following section gives some examples.

COVARIANCES AND SIMULATIONS

We have done several covariance analyses and simulations to see whether the data available under the assumptions stated above can achieve the accuracy required by AFF. We contemplate a scenario in which both range and phase measurements are available, as outlined above. The phases may be differenced as needed, and it may be necessary to initialize them by determining certain phase offsets. Issues to be addressed include:

- (1) What accuracy can be achieved with range measurements alone?
- (2) How can the phase data be used to best advantage? For example, what is the best differencing scheme?
- (3) Where should the transmitting and receiving antennas be located on the spacecraft?

- (4) For measurements that span several data epochs, how much spacecraft rotation (if any) is required in order to achieve sufficient accuracy? How long will the process take?
- (5) How many data epochs will it take to determine whatever phase constants are needed?
- (6) Are there any special problems, such as ambiguous solutions or poorly determined or highly correlated parameters?

The sections that follow address to some extent all of these questions except for number 3. First we present results obtained using software specially designed for the AFF problem. This software works in batch mode. That is, it supposes that enough epochs have elapsed so that a complete solution is specified and estimates the parameters for all the epochs in a single fit. A later section gives the results of simulations carried out using an extended version of the Real-Time GIPSY (GPS Inferred Positioning SYstem) software. This software works sequentially, as an operational system probably would, using whatever data are currently available to build a solution that becomes more precise as more epochs are added. To the extent that the two approaches are comparable, they reach consistent conclusions concerning the results obtainable from the postulated data.

Batch-Mode Covariance Analyses and Simulations

For the purposes of this report, we assume a realistic set of locations for the transmitting antenna and the three receiving antennas, common to all three spacecraft, as shown in Figure 1, where the antennas are at the four corners of the 1-meter-square front face of each spacecraft. Furthermore, a standard array configuration was used for all single-epoch analyses, an equilateral triangle with sides 1 km long and each spacecraft facing the center of the array. For multi-epoch measurements the spacecraft may have additional rotations, but their mean orientations will still be those of the standard configuration.

Range-Only Data, Single Epoch

The simplest covariance uses only the range data at a single epoch. In this scenario there are 18 measurements and 14 array parameters. For the standard array configuration, the covariance is well behaved: the correlation coefficient with the largest magnitude is -0.7319 . As a result, the standard deviations of the parameters are all consistent in a general way with the measurement errors of 1 cm. Because of the symmetry of the array, and the fact that the rotation angles are defined with respect to the rotated local coordinate frames, the uncertainties of corresponding orientation angles are the same for all three spacecraft: 0.0115 radians for ψ and θ , and 0.0153 radians for ϕ . From these results it is clear that even though range measurements are not in themselves capable of meeting the AFF accuracy requirements (particularly on the orientation angles), they are good enough to provide an excellent starting point for the phase solutions.

Ranges and Double-Differenced Phases, Two Epochs.

As a first try at introducing phase measurements into the solutions, we used data double-differenced between adjacent epochs and between receiving antennas on the same spacecraft. The rationale for this approach is that differencing between epochs removes the effect of phase offsets from the data, while differencing between receiving antennas focuses the observables on the poorly determined (relative to the requirements) orientation

parameters. The problem with this scheme is that while the change in orientation of the spacecraft (if any) between the first and second epochs is well determined, the orientations themselves are not. In effect, the phase observables act as linear constraints of the ranges: they link the two epochs, but the accuracy of the estimates still depends basically on the range measurements. For example, in a trial in which the nominal array parameters and orientations were used at both epochs, the uncertainties of the five non-angular parameters decreased on the average by 22%, while those of the nine angular parameters decreased by 30%. On the whole, the effect was similar to that of averaging the two sets of measurements.

Ranges and Single-Differenced Phases, Four Epochs

A more effective way to incorporate the phase data into the solution is to use single differences between receiving antennas on the same spacecraft. There are 12 algebraically independent differences of this kind at each epoch. On the other hand, without the time difference there are now 12 phase offsets (one for each difference) to be determined. If we ignore the range data for the moment and consider a phase-only solution for the orientation angles alone, then there are at each epoch 12 measurements, 9 orientation-angle parameters, and 12 phase-offset nuisance parameters. Hence the number of measurements equals or exceeds the number of parameters, and a solution is possible, when there are four or more epochs of data.

Another difference between this single-differencing approach and the double-differencing strategy described above is that the current method requires the orientation of the spacecraft to change between epochs. In fact, the uncertainties of the estimated angular parameters vary inversely as the square of the angular variation over a wide range. Figure 2 illustrates this relation for a specific case. A series of covariance calculations was done for the standard array configuration. A full set of angular offsets (three per spacecraft) was then selected for each of four epochs, by choosing values randomly and uniformly distributed over the range $(-0.40, +0.40)$ radians. In the series of trials these offsets were scaled by factors ranging from 0.1 to 3.0 and added to the basic angles. The covariance calculation was then performed for the four-epoch phase-only scenario outlined in the preceding paragraph. Examining the results, we separated the 12 phase-offset parameters from the 36 orientation-angle parameters and calculated the rms uncertainty of each group, with the results shown in the figure. Note that although the numerical values are nearly the same, the units of the phase-offset uncertainties are meters, and the units of the orientation-angle uncertainties are radians. The figure also includes a reference line with a slope of -2 ; it is evident that the inverse square law approximates the data well over a wide range of angular variation. However, as the maximum angular offsets increase to about ± 1 radian, the errors in the orientation angles predictably reach a minimum and begin to increase again. Incidentally, the error levels shown in the plot, unlike the slope, are of little significance, since widely variable results are obtained when the offset angles are chosen at random.

Since the use of single-differenced data requires spacecraft rotation, it is worthwhile to consider the time needed to rotate the spacecraft by the required amounts. For example, suppose that the array is operating and stationary when for some reason the AFF system is interrupted. How long would it take to rotate each spacecraft 90° about some axis and return to the observing configuration? If we assume that the 250-kg mass of the combiner spacecraft (the heaviest of the three) is uniformly distributed through its volume, and that the shape is a cube with a side of 1 meter, then the combiner's moment of inertia about each of the principal axes through the center of mass is

$$I = Ma^2/6 = 41.667 \text{ kg m}^2, \quad (3)$$

where M is the mass of the spacecraft and a is the length of a side.




Figure 2. Dependence of parameter errors on scale of spacecraft rotation

Furthermore, the available thrust from the pulsed plasma thrusters is $F = 4.5 \text{ mN}$, so the torque on one of the principal axes is just

$$\tau = Fa/2 = 2.25 \times 10^{-3} \text{ kg m}^2 \text{ s}^{-2}. \quad (4)$$

It is easy then to calculate that the time needed to accelerate from a standing start through an angle $\alpha/2$, then decelerate to a stop at angle α , and finally return to the starting orientation in the same way, is

$$t = 4 \left(\frac{\alpha I}{\tau} \right)^{1/2}. \quad (5)$$

For $\alpha = \pi/2$ and the values given above, $t = 682.2$ seconds. Evidently rotations happen slowly on these spacecraft. With a nominal AFF data rate of one per second, there will be more than enough time during a controlled rotation to collect the needed single-differenced data.

Since single differencing AFF phase data appears to be a workable way of reaching DS-3's angular accuracy requirement, we have carried the analysis further by running a

simulation of a four-epoch solution in which both range and single-differenced phase data are used to estimate all of the array parameters. At each epoch there are 30 observables, consisting of 18 ranges and 12 single-differenced phases, and the usual 14 array parameters. In addition there are 12 phase offsets that apply to all four epochs.

To set up the simulation, one first specifies the array configuration at each epoch, including three position parameters, two relative clock offsets, and nine orientation parameters. As discussed above, the orientations must be different at the different epochs, but the other parameters are generally kept the same. It is also necessary to specify the locations of the antennas in the local frame of the spacecraft. In the simulation described below we used the standard array configuration and added the same angular offsets used for the covariance calculations, distributed uniformly over the range $(-0.4, +0.4)$ radians.

The simulator then calculates the observables and adds Gaussian noise as specified by the user (in this case 1 cm on ranges and 10 μm on phases). When the observables are known, the solution proceeds in three steps.

First a relatively crude solution is done, one epoch at a time, using only the range data and assuming no a priori knowledge of the parameters to be estimated. The purpose of this step is simply to find a workable starting point for the least-squares refinements to follow.

In the second step of the solution, the crude range solution at each epoch is refined by means of a conventional iterative least-squares procedure. In the cases run to date, the "crude" solution is actually fairly good, with orientation angles within 0.1 radian of their refined values.

In step three the single-differenced phases are added to the input data, and all four epochs are combined in a larger least-squares calculation, with the refined range solution as a starting point. Now there are $4 \times (18 + 12) = 120$ measurements and $4 \times 14 + 12 = 68$ parameters. In our test case the range-plus-phase solution converged cleanly after three or four iterations, with the results shown in Table 1 (Case 1) for four classes of parameters. With respect to the performance requirements, only the coordinate and angular parameters are significant. The coordinates are somewhat better determined than they were by range data alone, and they meet the requirement easily. The angles, though much improved, are still not quite at the required level of 2.9×10^{-4} radians. Again, it should be understood that these results depend on arbitrary assumptions about the measurement errors, the placement of the antennas, and the orientation and rotation of the spacecraft. Actual performance will depend on the extent to which the system design and operating strategy can be improved.

An interesting feature of the parameter covariance for Case 1 is the fact that the angular errors are much larger (by a factor of 47) than the ratio of the measurement uncertainties (10^{-5} m) to the typical distance between receiving antennas (1 m). This result is not merely a geometric effect, but rather a reflection of the high correlation between the errors of the angular parameters and those of the phase offsets. To improve the angular errors further, it is necessary either to add measurements or to reduce the number of phase offsets.

In an ideal situation, one would design the signal paths from each of the three receiving antennas (on each spacecraft) to be exactly equal. In that case, all the single-differenced phase constants would be an integral number of cycles; and since the uncertainties of the phase offsets are a small fraction of a cycle, the integers could be

determined and the phase-offset parameters removed from the estimation process entirely. In practice, a design of the required precision (with phase offsets equal to 0.001 cycle or better at 30 GHz) is probably impractical. Alternatively, the differences between the phase paths might be calibrated before launch. In space, however, the calibration would change (as a result of differential heating as the spacecraft rotates, for example) at a significant level relative to the DS-3 requirements, on time scales of minutes and longer.

A more promising approach is to resolve double-differenced phases—that is, single-differenced phases differenced again between transmitting spacecraft. Variations of the phase offsets in these observables will be much smaller than those in the single-differenced phases. Two such double differences can be formed at each spacecraft, so that fixing their values leaves only six phase offsets to be estimated. In Table 1, Case 2 shows the result of this approach for our example. As expected, the angular errors are helped most by the reduction in the number of phase offsets, decreasing by a factor of 2.3. Formally, the orientations now satisfy the DS-3 requirement of 2.9×10^{-4} radians.

To complete the analysis, we also considered the possibility that all the phase offsets could be determined exactly, with the results shown in the table as Case 3. Again, the accuracy of the coordinate and time-offset parameters, which depend mostly on the range measurements, is hardly affected; but now the improvement in the angular measurements is quite dramatic. Clearly, the whole problem of dealing with phase offsets and calibration needs further attention; the subject comes up again in the following section.

Table 1
RMS ERRORS OF PARAMETERS IN THE FOUR-EPOCH SOLUTIONS

PARAMETER CLASS	CASE 1	CASE 2	CASE 3
Coordinate Parameters x_B, x_O, y_C (12 in all)	4.728 mm	4.645 mm	4.571 mm
Angular parameters $\psi_{A,B,C}, \theta_{A,B,C}, \phi_{A,B,C}$ (36 in all)	46.90×10^{-5} rad	20.57×10^{-5} rad	1.53×10^{-5} rad
Time-Offset parameters $R_{\tau AB}, R_{\tau AC}$ (8 in all)	3.347 mm	3.336 mm	3.333 mm
Phase-Offset parameters (Case 1: 12; Case 2: 6; Case 3: 0)	0.513 mm <i>or</i> 0.0513 cycle	0.238 mm <i>or</i> 0.0238 cycle	_____

Case 1: No cycle ambiguities fixed.

Case 2: 6 double-differenced cycle ambiguities fixed.

Case 3: All 12 single-differenced cycle ambiguities fixed.

Sequential-Mode Covariance Analyses and Simulations

In addition to the covariance analyses and simulations done in batch estimation mode as described above, we have also carried out simulations in sequential mode. JPL's Real-Time GIPSY (RTG) software has been extended to process the measurements, updating the spacecraft position and attitude parameters at each epoch as it would in the operational environment. The RTG estimation strategy also differs significantly from that used for batch

estimation. Here, the algorithm processes undifferenced pseudorange and phase measurements. For each combination of transmitting and receiving spacecraft, one time-varying phase bias (including clock errors) is estimated, together with two constant single-differenced (between receiving antennas) phase biases. The resolution of ambiguities is considered only for the differential phase biases. As in the batch-mode simulations, there are a total of 12 of these parameters.

This section differs from the previous one also in its approach to resolution of “integer-cycle” biases on the single-differenced phases. Here we suppose that as a result of careful design and calibration, the non-integral part of the single-differenced phase biases is a small fraction of an r-f cycle. In that case it is possible to determine the integral part of the bias and then solve for the small residual bias with a tight constraint on its magnitude. Furthermore, the small residual bias is assumed to arise primarily within the receiving spacecraft, so that it is the same for the signals from both transmitting spacecraft. As will be seen, this strategy improves the accuracy of the estimated orientation angles enough to meet the DS-3 requirement of 1 arcminute.

After generation of the simulated data, processing proceeds in two stages. In the first stage, pseudorange and carrier phase measurements are used to determine the spacecraft positions and to resolve the integral part of the single-differenced phase ambiguities. In the second, the phase data are used to refine the residual phase corrections and the spacecraft attitudes, while the phase and range data work together to improve the estimates of the position parameters and the clock offsets.

Simulation of Measurements

The simulated measurements are generated by first creating a time series of “truth” spacecraft attitudes and inter-spacecraft ranges. The error-free ranges and phases between receivers and transmitters are then calculated from these attitudes and ranges. Finally, clock errors, phase biases, and data noise are added to obtain simulated measurements. In this analysis, twenty epochs of data are simulated. The procedure for the simulation is as follows:

- 1) Specify the nominal array configuration. The spacecraft orientation angles and the locations of the antennas are those of the standard array configuration described above. The array geometry, however, is now slightly asymmetric and is parameterized directly in terms of the distances between the spacecraft reference points. In the obvious notation, $R_{AB} = 1000$ m, $R_{AC} = 1050$ m, and $R_{BC} = 1020$ m.
- 2) Add the desired attitude and range changes to the nominal values at each time epoch to create the truth time series.
- 3) Calculate the geometric ranges between receiver-transmitter pairs of different spacecraft at each epoch.
- 4) Add clock errors and phase biases to the observables at each epoch. These include:
 - a) for pseudorange measurements, clock errors for both the transmitter and the receiver;
 - b) for carrier phase measurements, clock errors and phase biases for both the transmitter and the receiver.
- 5) Add 1 cm of white noise to the pseudorange measurements and 10 μ m of white noise to the phase measurements. These are the same assumptions about data noise used in the batch simulations.

Results for Resolution of Integer-Cycle Phase Biases

As previously mentioned, the spacecraft need to rotate between epochs in this measurement scheme. In addition, the carrier phase measurement model needs reasonably good a priori knowledge of these rotations. As in the batch simulations, our approach here uses the attitude solution from pseudorange as the a priori information in the carrier phase measurement model. Although this solution is not accurate enough (about 0.4 degrees) to resolve phase ambiguities, it is adequate to provide nominal values to the phase measurement model.

In the estimation of both the spacecraft position and the attitude angles using the pseudorange measurements, the parameterization here is nearly the same as that used earlier. There are at each epoch a total of 18 measurements and 14 unknowns: 3 distances, 2 clock offsets, and 9 attitude angles. For the carrier phase measurements, however, the parameterization is a little different than it was in the batch estimation case, since the primary observables are now undifferenced. For each spacecraft at a single epoch, there are 9 unknowns: 3 attitude angles, 2 time-varying phase biases (including clock errors), and 4 constant single-differenced phase biases between receiving antennas; and there are 6 undifferenced measurements. After n epochs the number of parameters per spacecraft is therefore $(3n+2n+4)$. Thus, after four epochs the number of measurements equals the number of unknowns, and the phase data alone suffice to specify an attitude solution.

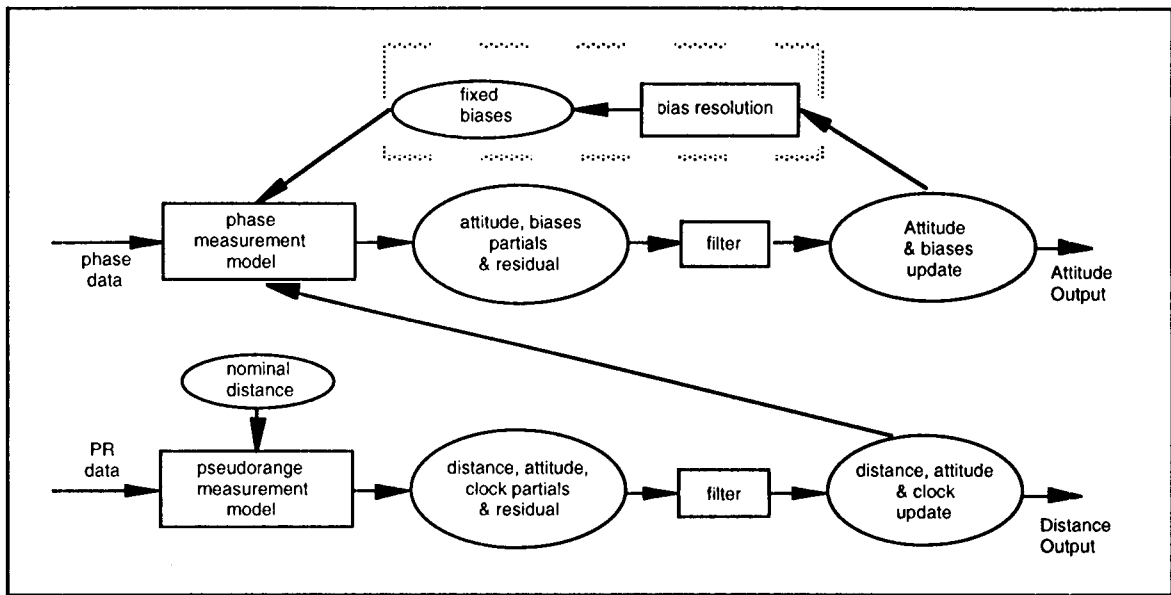


Figure 3. Data flow in AFF spacecraft relative position and attitude determination

Figure 3 shows the data flow in the filter for each epoch. First the pseudorange measurements are processed. The resulting attitude solution then supplies a priori values for updating the attitude and phase bias solutions with carrier phase measurements. This process continues at successive epochs until the solutions for the differential phase biases are accurate enough for ambiguity resolution. The ambiguity resolution technique itself is not included in this simulation. Our focus here is on how well this estimation technique can separate the differenced phase biases from the attitude parameters.




Figure 4. Formal errors of single-differenced phase bias parameters. The nomenclature specifying the differenced antennas is defined in the text.

In the particular simulation described here, the measurements are generated by adding an increment of 1 degree per epoch to all nine attitude angles to create the truth time series described above. A total of 20 epochs of measurements are simulated.

Since ambiguity resolution can be performed only on differenced phase biases, we examine the solution error and the covariance only for these parameters. As in the batch-mode simulations, there are four differenced phase ambiguities for each receiving spacecraft. Figure 4 shows the formal errors for the four differenced phase bias parameters for spacecraft A as a function of the number of epochs of data.

From these results we can see that a few degrees of attitude change makes it possible to separate the phase bias parameters from the attitude angles. After four to ten epochs the formal errors of the differenced phase biases drop to a couple of millimeters, and the actual errors in the differenced phase biases are generally less than a quarter wavelength, so that determination of the integral part of the biases is possible.

Results Using Carrier Phases with Integral-Cycle Ambiguities Resolved

Once the integral part of the single-differenced phase ambiguities has been resolved, the carrier phase becomes a stronger data type for attitude determination. At each epoch, the phases of all three receivers observing the same transmitter on another spacecraft now share a common bias, which remains unknown at this point. Since each spacecraft receives signals from two others, there are at each epoch 6 new phase measurements and 5 new parameters determined mostly by those data (2 unconstrained phase offsets and 3 orientation angles). There are also 2 small ($\ll 1$ cycle), tightly constrained residual single-differenced phase biases

for each receiving spacecraft, as mentioned in the previous section, that are postulated to remain effectively constant for many epochs.

To determine these parameters, RTG carries out a sequential estimation process involving four independent filtering steps. First, the 18 pseudorange measurements are used to update the 3 inter-spacecraft distances, the 2 clock biases, and the 9 attitude angles. Then the 6 phase measurements at each receiving spacecraft are used to refine the 3 attitude angles, the 2 common phase biases, and the 2 residual single-differenced phase biases, one spacecraft at a time. These last three filtering steps are carried out sequentially. Because the a priori values of the parameters are determined in the previous stage to well within the linear regime, no iterations are needed in this stage. Moreover, since the number of parameters estimated at each filtering step is small (between 7 and 14), the filter is small and fast—both desirable features for onboard autonomous estimation.




Figure 5. Spacecraft attitude angle errors when adjusting residual differential phase biases at three different levels of constraint

For a single epoch of data, the formal errors of the attitude angles depend strongly on the a priori constraint on the residual single-differenced phase biases. Figure 5 demonstrates this dependence for three levels of constraint. As usual, the measurement errors are assumed to be 1 cm for ranges and 10 μ m for phases. The figure shows that in this scenario the goal of 1-arcminute accuracy can be met only when the biases are smaller than about 0.5 mm.

If the constraints on the residual phase offsets are made infinitely tight (so that the single-differenced phase offsets are exactly integral), then the accuracy of the angular estimates improves further. In our example, the formal error on the inter-spacecraft distances is 0.41 cm, and the error on the clock offsets is 0.33 cm. The formal errors on the attitude angles range between 2.4 and 5.0 arc seconds. These results are quite similar to those for the analogous Case 3 in Table 1.

Comparison and Summary of the Batch and Sequential Analyses

In general, the good agreement between the error estimates calculated by the batch- and sequential-mode analyses confirms that the two independent calculations, approaching the problem from different points of view, are implemented correctly and giving credible results. Actually, the agreement is surprisingly good in light of the differences between the two approaches:

- 1) The batch simulation used somewhat larger spacecraft rotations, a factor to which the solution is quite sensitive (see Figure 2).
- 2) The solution is also quite sensitive to the specific rotations selected, particularly when the rotations are chosen at random, as was done in the batch-mode simulation.
- 3) The simulations used different numbers of epochs (20 and 1 in the sequential mode, 4 in the batch mode).
- 4) The parameterization of the constraints on the phase measurements was different.

The basic message from both analyses is that given the postulated measurement errors, the proposed implementation of AFF can attain the accuracy required by DS-3. Achieving those measurement errors will be a challenging technical problem, however. The following section touches on some of the relevant issues.

AREAS FOR FURTHER INVESTIGATION

Certain aspects of the AFF technique, both in theory and in application, have been given short shrift or none at all in the discussion above. The following is a short summary of the most important of these considerations, indicating in general terms how each can be handled, or what problems remain to be solved.

Initialization of the Array

At least as it applies to DS-3, AFF involves two kinds of initialization. The one already discussed applies when the relative locations and orientations of the spacecraft are already known at the 100-meter and 0.2-radian level. The array would need this kind of initialization during normal operation if communication among the spacecraft was briefly interrupted, for example. But when the three spacecraft are first deployed, and their orientation is virtually unknown, they need a more basic initialization; and the same situation might apply later, after a more serious anomaly in array communication or control. In these cases the main antennas might be pointing away from each other or rotating in an unknown way. An obvious solution is to add an auxiliary transmitting antenna and receiving antenna with their beams directed opposite to the main antennas. These auxiliary antennas would be used only for initialization, and the transmitter would broadcast a distinctive signal, using a different code or carrier frequency, or both, than the main transmitter. How such a scheme would be worked out obviously depends on the beam sizes of the various antennas, as discussed briefly below.

Locations of Antennas

Concerning the main antennas (one to transmit and three to receive) on each spacecraft, it seems obvious that they should be placed as far apart as possible, both to increase their angular discrimination and to reduce leakage of the transmitter signal into the

receivers. Thus, for purposes of the covariances calculated here, we have placed them at the four corners of the side of each spacecraft that faces the center of the array. Another possibility is to mount them on short masts that would allow them to be even farther apart.

Beam Shape of Antennas

A narrow beam (with high gain) has the advantage that it minimizes the required transmitter power and maximizes the ratio of received signal power (from another spacecraft) to leaked interfering power from a spacecraft's own transmitter. A wide beam has the virtue of straightforward recovery from large pointing errors. As a practical matter, each transmitter has to illuminate both receiving spacecraft, and each receiver has to see both transmitting spacecraft in the normal operating configuration. On balance, a circular beam with roughly constant gain out to 45 or 50 degrees off axis and a sharp drop-off at larger angles is probably the best choice. There seems to be little virtue in more complicated shapes.

Characterization and Control of Multipath

AFF's stringent accuracy requirements, particularly on the phase measurements, dictate that multipath be reduced to a low level and that the residual effect be calibrated to the extent possible. In practical terms that means careful attention to the mounting of the antennas and to the materials and geometry of the surface of the spacecraft. The fact that the operating geometry of the array is fixed means that the highest-accuracy measurements of multipath and of the effective locations of the antenna phase centers are needed only over a small range of directions. The immediate need is for some analytical work to see how the requirements can be met.

Ability to Transmit and Receive Simultaneously

According to the current AFF concept for DS-3, all three spacecraft will transmit and receive simultaneously and continuously. Inevitably, some of the transmitted signal at each spacecraft will leak into the front ends of the receivers, both by way of the transmitting and receiving antennas and by more devious routes. Without remediation this "self-signal" could easily saturate the front end of the AFF receivers or otherwise overwhelm the external signals. For example, under reasonable assumptions about the gains of the antennas and the isolation of the receivers, the ratio of powers entering the receiver from the local transmitter and a remote transmitter can easily be 45 dB or more. The potential for trouble is obvious.

Several remedial measures suggest themselves. If the transmitters operate at different frequencies, separated by something like 2 GHz, then an appropriate filter can reject most of the self signal, though at the cost of some fluctuation in the receiver delay. Another possibility is active rejection, in which the received signal is correlated against a replica of the self signal. The result is then used to control the amplitude and phase of an "anti-self" component, added to the received signal, that effectively cancels the unwanted input. This approach would complicate the receivers but could be very effective. If all else fails, transmission and reception can be time multiplexed, with the spacecraft taking turns not transmitting for an interval of perhaps a millisecond. This approach would certainly remove the original problem, but at considerable cost. It would reduce the total AFF data rate by a factor of three or more; it would require synchronization among the spacecraft; the non-simultaneous data would complicate processing; and transients associated with the cycling of the transmitters could cause problems. All these options are under consideration.

Effect on Observables of Instrumental Delays and Phase Shifts

Both range and phase observables are affected by instrumental delays, and phases are subject to other offsets as well. In the case of range data, the analysis presented above deals with these delays by applying two single-differenced clock offsets in the delay model. That is, it assumes that a single delay parameter characterizes the transmitter and all three receiver channels on each spacecraft. Careful calibration might justify this assumption at a particular epoch, but as time goes on differential effects will eventually (perhaps in a matter of minutes) invalidate the calibration and affect the observables at a significant level.

Phase data have similar problems, but the effect is somewhat different because the data are used differently (primarily to determine rotations rather than ranges). For phases, differencing is a useful tool that can reduce or simplify, but not remove, the effect of instrumental offsets on the observables. For undifferenced data there is an arbitrary phase offset. If phases are differenced between receivers on the same spacecraft (as was done above), careful design and calibration can bring that offset near an integral number of cycles. If the phases are double-differenced between transmitters and receivers, but not between epochs, the observable can be brought even closer to an integral number of cycles. If phases are differenced between epochs, a constant phase offset is suppressed completely. In fact, however, the phase shifts vary in time just as the instrumental delays do. To handle the data correctly, we need to know how these instrumental effects vary statistically as a function of time, and our knowledge of this behavior is rather poor at present.

Optimum Motion of Spacecraft While Determining Phase Offsets

In the batch-mode covariances and simulations, the rotations of the spacecraft between epochs were chosen at random. The reason for this apparently whimsical approach is that when rotations were chosen in a simple systematic way (such as uniform rotation about a single axis) the resulting covariances were generally very bad. Why the covariances behave in this way is not entirely clear; but it turned out that even when the rotations were chosen randomly, some choices led to anomalously large estimation errors, while a few gave anomalously small errors. Obviously we will benefit by taking the time to understand what kinds of motion optimizes the estimation process.

SUMMARY

Autonomous Formation Flying (AFF) is a spacecraft control technique applicable not only to DS-3 but to a variety of multi-spacecraft missions. In many cases, including DS-3, some form of AFF is essential to the accomplishment of the mission's objectives. This report has outlined a proposed GPS-like implementation of AFF for DS-3 and has provided a framework for predicting the performance of various measurement strategies.

Using this framework along with a realistic model of spacecraft geometry, we have performed covariance analyses which show that a suitable combination of range and phase measurements, with uncertainties of 1 cm and 10 μ m, respectively, can specify the locations and orientations of the three spacecraft at the required level of 1 cm for coordinates and 1 minute of arc for orientation angles. We have also implemented algorithms that perform a complete AFF solution of this kind on simulated data.

However, many problems remain to be solved before an AFF system can be implemented. Among the most crucial questions are how AFF can be initialized at the time

of deployment; how systematic measurement errors can be controlled to allow the required measurement accuracy; and how the effect of a spacecraft's transmitter on its received signals can be reduced to a manageable level.

ACKNOWLEDGMENT

The work described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.